



HAL
open science

Active Prediction for Discrete Event Systems

Stefan Haar, Serge Haddad, Stefan Schwon, Lina Ye

► **To cite this version:**

Stefan Haar, Serge Haddad, Stefan Schwon, Lina Ye. Active Prediction for Discrete Event Systems. 2020. hal-02951944v1

HAL Id: hal-02951944

<https://hal.science/hal-02951944v1>

Preprint submitted on 29 Sep 2020 (v1), last revised 11 Dec 2020 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Active Prediction for Discrete Event Systems

Stefan Haar¹, Serge Haddad², Stefan Schwoon², and Lina Ye³

¹ INRIA, LSV, ENS Paris-Saclay, CNRS, Université Paris-Saclay, France

² LSV, ENS Paris-Saclay, CNRS, INRIA, Université Paris-Saclay, France

³ CentraleSupélec, LRI, Université Paris-Saclay, France

Abstract. A central task in partially observed controllable system is to detect or prevent the occurrence of certain events called *faults*. Systems for which one can design a controller avoiding the faults are called *actively safe*. Otherwise, one may require that a fault is eventually detected, which is the task of *diagnosis*. Systems for which one can design a controller detecting the faults are called *actively diagnosable*. An intermediate requirement is *prediction*, which consists in determining that a fault will occur whatever the future behaviour of the system. When a system is not predictable, one may be interested in designing a controller to make it so. Here we study the latter problem, called *active prediction*, and its associated property, *active predictability*. In other words, we investigate how to determine whether or not a system enjoys the active predictability property, i.e., there exists an active predictor for the system.

Our contributions are threefold. From a semantical point of view, we refine the notion of predictability by adding two quantitative requirements: the minimal and maximal delay before the occurrence of the fault, and we characterize the requirements fulfilled by a controller that performs predictions. Then we show that active predictability is EXPTIME-complete where the upper bound is obtained via a game-based approach. Finally we establish that active predictability is equivalent to active safety when the maximal delay is beyond a threshold depending on the size of the system, and we show that this threshold is accurate by exhibiting a family of systems fulfilling active predictability but not active safety.

1 Introduction

Monitoring faulty systems. In monitoring faulty systems, two central tasks consist in detecting a fault that has occurred, resp. will occur, i.e. the tasks of *diagnosis* and *prediction*, respectively, based on observations. However, such tasks may be defeasible due to ambiguity (i.e. observations associated with both correct and faulty runs). In this case, one may introduce a *controller* to restrict the behaviour in order to enforce diagnosis (resp. prediction) to be processed. Such a controller is called an *active diagnoser* (resp. *active predictor*). Here we focus on the existence of an active predictor, a problem called *active predictability*.

Diagnosis. In partially observed discrete-event systems, diagnosis was defined and studied in the seminal paper by Sampath et al [16] (see also [6,7]). That work builds a deterministic version of the original model, a so-called *diagnoser*, that tries to detect the occurrence of faults. A system is called *diagnosable* if the diagnoser can detect every fault occurrence, possibly after some delay. As an illustration, consider the system in Figure 1, which we shall use as a running example, sometimes with different values for Σ_1 and Σ_2 , where Σ_1 and Σ_2 are subsets of events in the system. Precisely, $\Sigma_1, \Sigma_2 \subseteq \{a, b, c, d\}$, all of which are observable, while f represents a fault that is not directly observable. If, e.g., a is contained in both Σ_1 and Σ_2 , then the system is not diagnosable because any observation ada^n may belong to a faulty run or a correct one.

The diagnosability problem is in PTIME [21], via an approach called *twin-plant construction*. When the system is not diagnosable, it may have to be redesigned, e.g. by adding further sensors to enhance observability, or by forbidding some actions. Sampath et al [15] followed the last approach, called *active diagnosis*: one strives to synthesise a controller, based on partial observations, that forces the behaviour of the system to stay within a diagnosable subset of its behaviours. For instance, if the system in Figure 1 has $\Sigma_1 = \Sigma_2 = \{b\}$ and the controller has the right to block a , then the system is actively diagnosable.

The algorithm for the active-diagnosability problem in [15] operates in doubly exponential time. In [13], we revisited the problem using automata and game theory and established that in fact the active-diagnosability

problem is EXPTIME-complete. Later on, we generalised the framework, e.g. allowing the controller to be aware of deadlocks [4]. We also studied active diagnosis for probabilistic systems [1].

In loosely related works, Chantry and Pencolé [9] proposed a planning-based approach that allows the verdict of the diagnoser to be ambiguous; the works in [8,10,19] studied the problem of dynamic sensor activation to ensure some observation properties. In work more closely related to ours [18], Yin and Lafortune proposed a uniform approach for synthesizing property-enforcing supervisor by mapping the considered property to a suitably-defined information state, which is applicable to a class of properties that can be expressed in terms of such information states, including safety, diagnosability, opacity and so on. Note that predictability cannot be formulated as an information state in that framework since it depends also on future behaviours of the system; its enforcement thus requires new methods.




Fig. 1: Running example, with unobservable events indicated by dashed lines.

Prediction. Several works have studied the (passive) predictability problem, i.e. where no control is involved. For instance, if $\Sigma_1 = \{b\}$ and $\Sigma_2 = \{c\}$ in Figure 1, then upon first seeing c , an observer can predict that a fault will necessarily occur. In [11], Genc and Lafortune introduced a diagnoser construction to derive a necessary and sufficient condition for predictability in systems modeled by regular languages. Ye, Dague, and Nouioua [17] proposed a polynomial time algorithm for predictability analysis in a centralized way and then extend it to a distributed framework. Brandon Briones and Madalinski [5] introduced and studied two variants of predictability by defining an additional requirement about either a lower bound or an upper bound on the number of events between the fault prediction and the fault occurrence. Then Yin and Li [20] investigated the bounded predictability in the decentralized framework, and proposed a polynomial-time algorithm for its verification. Madalinski and Khomenko [14] reduce the predictability problem for a Petri net to LTL-X model checking. All these previous works focus on passive predictability.

Our contributions. First we refine the paradigm of prediction by allowing an observer to quantify its observations. Unlike [5] but similar to [20], our predictors can at the same time provide both lower and upper bounds on the number of observations before a fault may (resp. must) occur. For instance, upon seeing c in the previous example, an observer can not only predict that a fault will eventually happen but that it will necessarily happen between the first and the second observable event after c . In practice, if a fault prediction is issued, the reaction procedure of the system should be triggered. As such interventions may require a certain amount of time to take effect, having both lower and upper bounds are relevant performance criteria for capture such timing issues.

We then turn to the case of active prediction, where a controller tries to restrict the system's behaviour so that faults can be reliably predicted. For instance, if $\Sigma_1 = \{a, b\}$ and $\Sigma_2 = \{a, c\}$ in Figure 1, then faults are unpredictable, but if a controller has the right to block a , it becomes actively predictable (with the aforementioned bounds). We formalize the idea of active predictability and then propose a class of controller, called active predictor. We then show that active predictability is equivalent to the existence of an active predictor.

Next, we focus on the decision and synthesis problems, i.e. to decide whether the system is actively predictable, and if so, how to build an active predictor. In active *diagnosability* [13], the solution exploited the fact that whether a sequence of observations is ambiguous (i.e. corresponds to both faulty and correct runs) is *independent of the control* that was applied in the past. In prediction, by contrast, the eventuality of a fault occurrence in the future *depends on the control* that is going to be applied. Thus solving the active-predictability problems requires new techniques.

We establish that the decision problem is EXPTIME-complete by reducing it to a turn-based game with a Büchi objective of exponential size. A memoryless winning strategy of this game provides the main ingredient

to build an active predictor. Furthermore we show that instead of solving this Büchi game (which takes quadratic time), one can equivalently in linear time (1) solve a reachability game, (2) build a safety game that depends on the winning states of the reachability game, and (3) solve it and combine the winning strategies to get a winning strategy for the Büchi game when it exists.

Finally we study the relation between the lower prediction bound k and the number of states n of the system. We establish that if $k \geq 2^n$ then a system is k -actively predictable if and only if it is actively safe. This bound is tight since we exhibit a family of systems of size $\mathcal{O}(n)$ such that the system is 2^n -actively predictable but not actively safe.

Organization. In Section 2, we introduce prediction in both the uncontrollable and controllable framework and establish a class of controller called *active predictor*. The existence of such a controller is equivalent to active predictability. The construction of an active predictor (if it exists) is carried out in Section 3, providing simultaneously the solutions to the decision and synthesis problems. Section 4 complements these results by a tight analysis of complexity bounds. We conclude and give some perspectives to this work in Section 5.

2 The Active Prediction Problem

As usual, for an alphabet Σ , we use Σ^* and Σ^ω , to denote the finite and infinite words over Σ , and $\Sigma^\infty := \Sigma^* \cup \Sigma^\omega$. The length of a word $\sigma \in \Sigma^*$ is denoted $|\sigma|$, and \preceq represents the prefix notation.

Labeled transition systems

When dealing with discrete event systems (DES), systems are often modeled using labeled transition systems (LTS).

Definition 1. A labeled transition system is a tuple $\mathcal{A} = \langle Q, q_0, \Sigma, T \rangle$ where:

- Q is a set of states with $q_0 \in Q$ the initial state;
- Σ is a finite set of events;
- $T \subseteq Q \times \Sigma \times Q$ is a set of transitions.

We note $q \xrightarrow{a}_{\mathcal{A}} q'$ for $(q, a, q') \in T$; this transition is said to be *enabled* in q . A *run* over the infinite word $\sigma = a_1 a_2 \dots \in \Sigma^\omega$ is a sequence of states $(q_i)_{i \geq 0}$ with $q_i \xrightarrow{a_{i+1}}_{\mathcal{A}} q_{i+1}$ for all $i \geq 0$, and we write $q_0 \xrightarrow{\sigma}_{\mathcal{A}}$ if such a run exists. A finite run over $\sigma \in \Sigma^*$ is defined analogously, and we write $q \xrightarrow{\sigma}_{\mathcal{A}} q'$ if it ends at state q' . A state q is *reachable* if there exists a run $q_0 \xrightarrow{\sigma}_{\mathcal{A}} q$ for some σ . The index \mathcal{A} in those relations will be omitted if unambiguous.

In order to formalize problems related to prediction, we partition Σ into two disjoint sets Σ_o and Σ_{uo} , the sets of *observable* and of *unobservable events*, respectively. Moreover, we distinguish a special *fault* event $f \in \Sigma_{uo}$. We say σ is *correct* if $\sigma \in (\Sigma \setminus \{f\})^*$ (we will denote $\Sigma \setminus \{f\}$ with the short form $\Sigma^{\setminus f}$ in the following), and that σ is *faulty* otherwise. For $\Sigma' \subseteq \Sigma$, define its projection $\mathcal{P}_{\Sigma'}(\sigma)$ inductively by: $\mathcal{P}_{\Sigma'}(\varepsilon) = \varepsilon$; $\mathcal{P}_{\Sigma'}(\sigma a) = \mathcal{P}_{\Sigma'}(\sigma) a$ when $a \in \Sigma'$, and $\mathcal{P}_{\Sigma'}(\sigma a) = \mathcal{P}_{\Sigma'}(\sigma)$ otherwise. For the sake of simplicity, write \mathcal{P} for \mathcal{P}_{Σ_o} , $|\sigma|_o$ for $|\mathcal{P}(\sigma)|$, $|\sigma|_{\Sigma'}$ for $|\mathcal{P}_{\Sigma'}(\sigma)|$, and for $a \in \Sigma$, write $|\sigma|_a$ for $|\sigma|_{\{a\}}$. When σ is an infinite word, its projection is the limit of the projections of its finite prefixes. This projection can be either finite or infinite. As usual the projection is extended to languages.

Definition 2 (Languages of an LTS). Let $\mathcal{A} = \langle Q, q_0, \Sigma, T \rangle$ be an LTS. The finite and the infinite (correct) languages of \mathcal{A} are defined by:

- $\mathcal{L}^*(\mathcal{A}) = \{ \sigma \in \Sigma^* \mid \exists q, q_0 \xrightarrow{\sigma}_{\mathcal{A}} q \}$ and $\mathcal{L}^\omega(\mathcal{A}) = \{ \sigma \in \Sigma^\omega \mid q_0 \xrightarrow{\sigma}_{\mathcal{A}} \}$;
- $\mathcal{L}_c^*(\mathcal{A}) = \{ \sigma \in (\Sigma^{\setminus f})^* \mid \exists q, q_0 \xrightarrow{\sigma}_{\mathcal{A}} q \}$ and $\mathcal{L}_c^\omega(\mathcal{A}) = \{ \sigma \in (\Sigma^{\setminus f})^\omega \mid q_0 \xrightarrow{\sigma}_{\mathcal{A}} \}$

\mathcal{A} is *safe* if $\mathcal{L}^*(\mathcal{A}) = \mathcal{L}_c^*(\mathcal{A})$ (i.e. no fault ever occurs).

The union of finite and infinite languages of \mathcal{A} is denoted $\mathcal{L}^\infty(\mathcal{A}) = \mathcal{L}^*(\mathcal{A}) \cup \mathcal{L}^\omega(\mathcal{A})$. The inverse observable projection with respect to \mathcal{A} and $w \in \Sigma_o^*$ is defined as $\mathcal{P}_\mathcal{A}^{-1}(w) = \{\sigma \in \mathcal{L}^*(\mathcal{A}) \mid \mathcal{P}(\sigma) = w\}$, which can be simply denoted by $\mathcal{P}^{-1}(w)$ if there is no ambiguity. An LTS \mathcal{A} is *deterministic* if for every pair $q \in Q, a \in \Sigma$ there is at most one q' such that $q \xrightarrow{a} q'$. For a deterministic LTS we write $T(q, a) = q'$ if $q \xrightarrow{a} q'$. As is the case for classical diagnosis problems, we make two **assumptions** on \mathcal{A} :

- Liveness: $\forall q \in Q, \exists a, q', q \xrightarrow{a} q'$.
- Convergence: $\mathcal{L}^\omega(\mathcal{A}) \cap \Sigma^* \Sigma_{uo}^\omega = \emptyset$.

Liveness implies that from any reachable state of an LTS, there exists at least one transition enabled in that state. Convergence guarantees that there is no infinite sequence of unobservable events. When \mathcal{A} is convergent, then for all $\sigma \in \mathcal{L}^\omega(\mathcal{A})$, one has $\mathcal{P}(\sigma) \in \Sigma_o^\omega$.

Example 3. Figure 1 shows a live and convergent LTS with $\Sigma_o = \{a, b, c, d\}$, $\Sigma_{uo} = \{f\}$, $\Sigma_1 \subseteq \Sigma_o$, $\Sigma_2 \subseteq \Sigma_o$ and $\Sigma_1 \cup \Sigma_2 \neq \emptyset$. Transitions labelled by unobservable events are dashed. We also factorize transitions with same source and target states. Depending on Σ_1 and Σ_2 , this LTS may have different levels of predictability (see Example 6 for further explanation).

Predictability

Intuitively, a system is predictable with respect to a fault f if in every faulty run, an observer can be certain that f is going to occur before it actually happens. Before formally defining predictability, we first present some useful notations. Given $\sigma \in \mathcal{L}^\infty(\mathcal{A})$ and $n \leq |\sigma|_o$, $pre_n(\sigma)$ denotes the minimal (w.r.t. \preceq) prefix of σ such that $|pre_n(\sigma)|_o = n$. As an abbreviation, $pre(\sigma) := pre_{|\sigma|_o}(\sigma)$ removes unobservable events from the end of σ . For example, in the LTS of Figure 1, we have (as f is unobservable) $pre_0(bdf) = \varepsilon$, $pre_1(bdf) = b$ and $pre(bdf) = pre_2(bdf) = bd$. We naturally extend pre to sets of words.

An observed sequence w forbids prediction of a fault when a correct, infinite future behavior is still possible. We introduce different kinds of observed sequences.

Definition 4. (*observation properties*) Let \mathcal{A} be an LTS, $w \in \Sigma_o^*$, and $m \in \mathbb{N}$. Then w is:

- surely correct in \mathcal{A} if $pre(\mathcal{P}_\mathcal{A}^{-1}(w)) \cap \Sigma^* f \Sigma^* = \emptyset$;
- surely faulty in \mathcal{A} if $\mathcal{P}_\mathcal{A}^{-1}(w) \cap \mathcal{L}_c^*(\mathcal{A}) = \emptyset$;
- ambiguous in \mathcal{A} if it is neither surely correct nor surely faulty in \mathcal{A} ;
- m -correct in \mathcal{A} if ww' is surely correct in \mathcal{A} for all $w' \in \Sigma_o^m$;
- m -faulty in \mathcal{A} if ww' is surely faulty in \mathcal{A} for all $w' \in \Sigma_o^m$;
- ω -faulty in \mathcal{A} if there exists $m \in \mathbb{N}$ such that w is m -faulty.

We now define the notion of k - l -predictability, which means that the occurrence of a fault can be predicted with certainty, based on what has been observed so far, at least k observations before it does occur, and such that the fault definitely occurs before the l -th additional observation. In the sequel, \mathbb{N}^+ denotes $\mathbb{N} \setminus \{0\}$ and \mathbb{N}_ω (resp. \mathbb{N}_ω^+) denotes \mathbb{N} (resp. \mathbb{N}^+) enlarged with ω which is an absorbing item for addition.

Definition 5. (*Predictability*) Let \mathcal{A} be an LTS, $w \in \Sigma_o^*$, $k \in \mathbb{N}$, and $l \in \mathbb{N}_\omega^+$.

- w is k - l -faulty in \mathcal{A} if w is k -correct and $(k+l)$ -faulty in \mathcal{A} .
- \mathcal{A} is k - l -predictable if for all $\sigma f \in \mathcal{L}^*(\mathcal{A})$, $\mathcal{P}(\sigma)$ has a k - l -faulty prefix.

Remark 1. If w is k - l -faulty in \mathcal{A} , then w is also k' - l' -faulty in \mathcal{A} for all $k' \leq k$ and $k' + l' \geq k + l$.

As an abbreviation, we will call \mathcal{A} k -predictable if it is k - ω -predictable, and simply *predictable* if it is 0-predictable. Thus, Remark 1 implies that predictability is weaker than any other notion of k - l -predictability.

Example 6. Consider the LTS of Figure 1:

- it is not predictable if $\Sigma_1 \cap \Sigma_2 \neq \emptyset$;

- it is 1-1-predictable and not 2-predictable if $\Sigma_1 \cap \Sigma_2 = \emptyset$, and both of them are not empty;
- it is 2-1-predictable if $\Sigma_1 = \emptyset$ and $\Sigma_2 \neq \emptyset$.

Let us focus on relations between these notions in finite LTS. Especially, Proposition 8 establishes bounds for predictability in finite LTS.

Lemma 7. *Let \mathcal{A} be an LTS with n states, where n is finite, and $w \in \Sigma_o^*$. Then w is ω -faulty if and only if it is n -faulty.*

Proof. One direction of the proof is trivial. For the other, suppose that w is not n -faulty. Then there exists $w' \in \Sigma_o^n$ such that ww' is not surely faulty, which entails that there exists some $\sigma' \in \mathcal{P}^{-1}(ww') \cap \mathcal{L}_c^*(\mathcal{A})$, i.e. σ' does not contain f . Consider some run of \mathcal{A} over σ' , and let q_i , for $i = 0, \dots, n$, be the state reached in that run after $pre_{|w|_o+i}(\sigma')$; hence q_0 is reached by a run over some word in $\mathcal{P}^{-1}(w)$. By the pigeonhole principle, there are two distinct indices $i < j$ with $q_i = q_j$, which allows to construct a word $\sigma'' \in \mathcal{L}_c^\omega(\mathcal{A})$ such that w is a prefix of $\mathcal{P}(\sigma'')$. Now \mathcal{A} is convergent, so $\mathcal{P}(\sigma'') = ww''$ for some $w'' \in \Sigma_o^\omega$ and ww'' is not surely faulty, hence w is not ω -faulty. \square

Proposition 8. *Let \mathcal{A} be a k -predictable LTS with n states, where n is finite.*

- (i) \mathcal{A} is k - n -predictable.
- (ii) If \mathcal{A} is not safe, then $k < n$.

Proof.

- (i) Suppose that \mathcal{A} is not k - n -predictable. Then there exists some $\sigma f \in \mathcal{L}^*(\mathcal{A})$ such that any k -correct prefix $w \preceq \mathcal{P}(\sigma)$ is not n -faulty. But since \mathcal{A} is k -predictable, there must be at least some $w \preceq \mathcal{P}(\sigma)$ that is both k -correct and ω -faulty, which is a contradiction by Lemma 7.
- (ii) Suppose that \mathcal{A} is k -predictable for $k \geq n$. If there is some $\sigma f \in \mathcal{L}^*(\mathcal{A})$, then since \mathcal{A} is k -predictable, there exists some k - ω -faulty prefix $w \preceq \mathcal{P}(\sigma)$. Since \mathcal{A} is live and convergent, there exists some $w' \in \Sigma_o^k$ such that $ww' \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}))$. Since w is k -correct, ww' is surely correct, hence w is not n -faulty. Then by Lemma 7, w cannot be ω -faulty, a contradiction. \square

Active predictability

We suppose that Σ_o is partitioned into subsets $\Sigma_c \subseteq \Sigma_o$ of *controllable* and $\Sigma_{uco} = \Sigma_o \setminus \Sigma_c$ of *uncontrollable* actions. Intuitively, a controller may forbid a subset of the controllable actions based on the observations made so far, thereby restricting the behaviour of \mathcal{A} .

Definition 9 (Controlled LTS). *Let \mathcal{A} be an LTS. A controller for \mathcal{A} is a mapping $cont : \mathcal{P}(\mathcal{L}^*(\mathcal{A})) \rightarrow 2^\Sigma$ such that for all w , $\Sigma_{uco} \cup \Sigma_{uo} \subseteq cont(w)$. The controlled LTS $\mathcal{A}_{cont} = \langle Q_{cont}, q_{0cont}, \Sigma, T_{cont} \rangle$ is defined as the smallest LTS satisfying:*

- $q_{0cont} = \langle \varepsilon, q_0 \rangle \in Q_{cont}$;
- if $\langle w, q \rangle \in Q_{cont}$, $a \in cont(w)$, and $q \xrightarrow{a}_{\mathcal{A}} q'$, then $\langle w\mathcal{P}(a), q' \rangle \in Q_{cont}$ and $\langle w, q \rangle \xrightarrow{a}_{\mathcal{A}_{cont}} \langle w\mathcal{P}(a), q' \rangle$.

The goal of our controllers is to make the system predictable by preserving liveness and to perform prediction at the same time. Before formally defining prediction verdicts in Definition 11, we discuss their intuitive meanings: \top means that the controller is currently unable to predict a fault, while $\langle k, l \rangle$ means that the run is correct so far but a fault can be predicted to happen between the next k and $k + l$ observations. When $l = \omega$, a fault is predicted but without an upper bound. $\langle ?, m \rangle$ means that a fault may or may not have happened yet but one will surely occur within m further observations, and \perp means that a fault has definitely already occurred.

Example 10. Consider again the LTS from Figure 1 and assume that $\Sigma_1 = \{a\}$ and $\Sigma_2 = \{b\}$. At the beginning, no fault can be predicted, so a controller would be expected to emit the prediction \top . After observing b , the controller could predict that a fault will happen between the first and second next observation to come, i.e. $\langle 1, 1 \rangle$. After seeing d , this would change to $\langle 0, 1 \rangle$, and finally to \perp .

Definition 11 (predictions). Let $\mathbb{P} := \{\top\} \cup (\mathbb{N} \times \mathbb{N}_\omega^+) \cup (\{?\} \times \mathbb{N}_\omega^+) \cup \{\perp\}$ be the set of possible predictions. We define the following measures $\kappa, \mu : \mathbb{P} \rightarrow \mathbb{N}_\omega \cup \{-1, \omega + 1\}$:

- $\kappa(\top) = \omega + 1$, $\kappa(\langle k, l \rangle) = k$, and $\kappa(p) = -1$ otherwise;
- $\mu(\top) = \omega + 1$, $\mu(\langle k, l \rangle) = k + l$, $\mu(\langle ?, m \rangle) = m$, and $\mu(\perp) = 0$.

We also define two particular types of subsets of \mathbb{P} : For $k \in \mathbb{N}$ and $l \in \mathbb{N}^+$, let $\mathbb{P}_{k,l} := \{\top, \perp\} \cup \{\langle k', l' \rangle \mid k' \leq k, l' \leq l\} \cup \{\langle ?, m \rangle \mid m < l\}$ and $\mathbb{P}_{k,\omega} := \{\top, \perp, \langle ?, \omega \rangle\} \cup \{\langle k', \omega \rangle \mid k' \leq k\}$.

The values $\kappa(p)$ and $\mu(p)$ define the ‘window’ (lower and upper bound on future observations) within which a fault is to occur according to prediction p . Here, -1 indicates that the fault may or must have occurred in the past, and in the case of \top , $\omega + 1$ is chosen for technical convenience. A predictor using values from $\mathbb{P}_{k,l}$ makes firm commitments on both the lower and upper bounds within which a fault is going to occur, while a predictor with values from $\mathbb{P}_{k,\omega}$ only commits to a lower bound.

Definition 12 (compatible predictions). Let $p, p' \in \mathbb{P}$ and $k \in \mathbb{N}, l \in \mathbb{N}_\omega^+$. We say that $\langle p, p' \rangle$ are k - l -compatible if the following conditions are all satisfied:

- if $p = \top$, then $\kappa(p') \geq k$ else $\kappa(p') \geq \kappa(p) - 1$;
- $\mu(p') \leq \mu(p)$, and if $0 < \mu(p) < \omega$, then $\mu(p') < \mu(p)$;
- if $p' \neq \top$, then $\mu(p') \leq k + l$.

Moreover, p is called k - l -initial if $\langle \top, p \rangle$ are k - l -compatible.

The conditions in Definition 12 describe the relations that should reasonably hold between a prediction p made for some observation w and the prediction p' made when one has observed one additional event in a k - l -predictable controlled LTS. Intuitively these are:

1. When a fault is first predicted, it should be at least k observations in advance, and the gap between this lower bound and the upper bound should be at most l . This is why $p = \top$ should imply $\kappa(p') \geq k$. In particular, one cannot switch from \top to $\langle k', l' \rangle$ for any $k' < k$, nor directly to $\langle ?, m \rangle$ or \perp . Moreover, the third condition ensures that when switching from \top to $\langle k', l' \rangle$, we have $k' + l' \leq k + l$, which with $k' \geq k$ implies $l' \leq l$.
2. Having predicted a fault within a certain ‘window’, the subsequent predictions can only become more precise. Thus, one can maintain or shrink that window, but not enlarge, shift, or forget about it. Figure 2 illustrates this idea. E.g., when a predictor announces a fault between the 3rd and 7th following observation, expressed by $p = \langle 3, 4 \rangle$, then one step later it must give $p' = \langle 2, 4 \rangle$ or a more precise verdict such as $\langle 3, 2 \rangle$. As another example, if the controller has arrived at a verdict of $\langle ?, 6 \rangle$, meaning “a fault has occurred, or will occur within six further observations”, then the information gained from an additional observation may lead it to conclude that the fault has now definitely occurred (\perp), will occur later (e.g., $\langle 1, 3 \rangle$), or to maintain the prediction (e.g., $\langle ?, 5 \rangle$). Note that $\langle ?, 6 \rangle$ could only be reached by passing through $\langle 0, m \rangle$, for some $m > 6$, earlier in the observation. These relations are ensured by allowing κ to decrease by at most one and requiring μ to strictly decrease (if an upper bound was given).

A k - l -initial prediction is one that is admissible for the empty observation.

Definition 13 (active predictor). Let \mathcal{A} be an LTS, $\mathbb{P}' \subseteq \mathbb{P}$, and $h = \langle \text{cont}, \text{pred} \rangle$, where cont is a controller and pred is a mapping from $\mathcal{P}(\mathcal{L}^*(\mathcal{A}_{\text{cont}}))$ to \mathbb{P}' . We call h a k - l -active predictor over \mathbb{P}' , for $k \in \mathbb{N}$ and $l \in \mathbb{N}_\omega^+$, if and only if:

- (i) $\mathcal{A}_{\text{cont}}$ is live;




Fig. 2: Examples of compatible predictions $\langle p, p' \rangle$ after n resp. $n + 1$ observations, where p is illustrated above the timeline, and p' is one of the predictions below. Solid intervals indicate periods in which a fault is predicted.

- (ii) $\text{pred}(\varepsilon)$ is k - l -initial;
- (iii) for $w \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{\text{cont}}))$, the prediction satisfies the following:
 - if $\text{pred}(w) = \top$, then w is $(k + 1)$ -correct in $\mathcal{A}_{\text{cont}}$;
 - if $\text{pred}(w) = \langle k', l' \rangle$, then w is k' - l' -faulty in $\mathcal{A}_{\text{cont}}$;
 - if $\text{pred}(w) = \langle ?, m \rangle$, then w is ambiguous and m -faulty in $\mathcal{A}_{\text{cont}}$;
 - if $\text{pred}(w) = \perp$, then w is surely faulty in $\mathcal{A}_{\text{cont}}$;
- (iv) for $a \in \Sigma_o$, $w, wa \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{\text{cont}}))$, $\langle \text{pred}(w), \text{pred}(wa) \rangle$ are k - l -compatible.

Intuitively, condition (ii) requires that the control cannot introduce deadlocks, and conditions ((ii)),((iii)) ensure that the predictions have the intended semantics. Condition ((iv)) ensures compatibility between two subsequent predictions along an observation. If there exists a k - l -active predictor for \mathcal{A} , we call \mathcal{A} k - l -active-predictable, or just *actively predictable*. Moreover, \mathcal{A} is called *actively safe* if there exists an active predictor for \mathcal{A} over $\{\top\}$, which entails that $\mathcal{A}_{\text{cont}}$ is safe.

Example 14. In the LTS \mathcal{A} of Figure 1, assume that $\Sigma_1 = \{a, c\}$, $\Sigma_2 = \{a, b\}$, $\Sigma_c = \{a, b, c\}$. Let $h = \langle \text{cont}, \text{pred} \rangle$ be defined by:

- $\text{cont}(\varepsilon) = \{b, c, d, f\}$, and $\text{cont}(w) = \Sigma$ otherwise;
- $\text{pred}(\varepsilon) = \text{pred}(w) = \top$, where $w \in c\Sigma_o^* \cap \mathcal{P}(\mathcal{L}^*(\mathcal{A}))$, $\text{pred}(b) = \langle 1, 1 \rangle$, $\text{pred}(bd) = \langle 0, 1 \rangle$, and $\text{pred}(bda^+) = \perp$.

In this example, h is a 1-1-active predictor.

Proposition 15 and Proposition 17 will exhibit a tight correspondence between the existence of a k - l -predictor for \mathcal{A} and the existence of a controller that makes \mathcal{A} k - l -predictable. Additionally, Proposition 17 shows that the set of predictions used in a predictor can be limited to a finite set, either committing the prediction to a lower and upper bound for the occurrence of a fault, or just a lower bound.

Proposition 15. *If $h = \langle \text{cont}, \text{pred} \rangle$ is a k - l -active predictor for an LTS \mathcal{A} , then $\mathcal{A}_{\text{cont}}$ is k - l -predictable.*

Proof. Let $\sigma f \in \mathcal{L}^*(\mathcal{A})$. We need to prove that $\mathcal{P}(\sigma)$ has a k - l -faulty prefix. Due to Definition 13((ii)) either $\text{pred}(\varepsilon) = \langle k', l' \rangle$, for some $k' \geq k$ and $k' + l' \leq k + l$, then by Definition 13((iii)) and Remark 1, ε is k - l -faulty, or $\text{pred}(\varepsilon) = \top$, so ε is $k + 1$ -correct. Then, since $\mathcal{P}(\sigma)$ is not 1-correct and due to Definition 13((iii)), $\text{pred}(\mathcal{P}(\sigma)) \neq \top$. So let wa be the minimal prefix of $\mathcal{P}(\sigma)$ such that $\text{pred}(wa) \neq \top$. Then $\text{pred}(w) = \top$, w is $k + 1$ -correct, and wa is still k -correct, therefore $\text{pred}(wa)$ is of the form $\langle k', l' \rangle$ with $l' \leq l$ because $\text{pred}(wa) \neq \top$ and all other cases are excluded by Definition 13((iii)). Furthermore, $\kappa(\text{pred}(w)) = \omega + 1$, $\kappa(\text{pred}(wa)) = k'$, so because of Definition 13((iv)), we have that $k' \geq k$. Thus, wa is k' - l' -faulty, hence $|\mathcal{P}(\sigma)| \geq |wa| + k'$. Consider the prefix w' of length $|wa| + k' - k$ of $\mathcal{P}(\sigma)$; w' must be k - l' -faulty, and hence with Remark 1 we get that w' is k - l -faulty. \square

Lemma 16. *Let \mathcal{A} be a k - l -predictable LTS and $w \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}))$. If w is k -correct but not $k + 1$ -correct, for $k \in \mathbb{N}$, then w is k - l' -faulty for some $l' \leq l$.*

Proof. Since w is k -correct but not $k+1$ -correct, there is some $\sigma \in \mathcal{P}^{-1}(w\Sigma_o^k)$ such that $\sigma f \in \mathcal{L}^*(\mathcal{A})$. Since \mathcal{A} is k - l -predictable, $\mathcal{P}(\sigma)$ has a k - l -faulty prefix $w' \preceq w$. If l is finite, then let $l' := l - |w| + |w'|$; otherwise, let $l' = m - k - |w| + |w'|$, where $m \in \mathbb{N}$ is a value such that $\mathcal{P}(\sigma)$ is m -faulty. In either case, $l' \leq l$. Since w' is $k+l$ -faulty, w must be $k+l'$ -faulty and hence also k - l' -faulty. \square

Proposition 17. *Let \mathcal{A} be an LTS. If there exists a controller $cont$ such that \mathcal{A}_{cont} is live and k - l -predictable, then there exist k - l -active predictors $h = \langle cont, pred \rangle$ for \mathcal{A} over both $\mathbb{P}_{k,l}$ and $\mathbb{P}_{k,\omega}$.*

Proof. We only need to construct a suitable function $pred$ that satisfies Definition 13 ((ii)), ((iii)), and ((iv)). We first construct $pred$ over \mathbb{P} in general and then discuss the restrictions to $\mathbb{P}_{k,l}$ and $\mathbb{P}_{k,\omega}$ (which are identical when $l = \omega$).

If \mathcal{A}_{cont} is k - l -predictable, then either ε is $k+1$ -correct and we set $pred(\varepsilon) := \top$, or otherwise ε is k - l' -faulty with $l' \leq l$ (by Lemma 16) and we set $pred(\varepsilon) := \langle k, l' \rangle$, thus ((ii)) is satisfied. As for ((iii)) and ((iv)), we proceed by induction on the length of $w \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$. So let $w, wa \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$, and assume that ((iii)) holds for w . In the following, we shall abbreviate $pred(w), pred(wa)$ by p, p' , respectively. We proceed by a case-by-case analysis on p and show that in each case one can find a value for p' that verifies ((iii)). The reader can verify that p and p' satisfy ((iv)) in every case. Moreover, the choice of p' maintains the invariant that if $p' = \langle k', l' \rangle$ or $p' = \langle ?, m \rangle$, then $k' \leq k$ and $l', m \leq l$.

- If $p = \top$, then w is $k+1$ -correct.
 - Either wa is $k+1$ -correct, then $p' := \top$.
 - Or wa is not $k+1$ -correct but k -correct, then by Lemma 16 it is k - l' -faulty for some $l' \leq l$ and $p' := \langle k, l' \rangle$.
- If $p = \langle k', l' \rangle$ with $k' > 0$ and $l' \leq l$, then w is k' - l' -faulty and wa is $(k' - 1)$ - l' -faulty, so $p' := \langle k' - 1, l' \rangle$.
- If $p = \langle 0, l' \rangle$ with $l' \leq l$, then w is l' -faulty.
 - Either wa is surely correct. Then wa is $(l' - 1)$ -faulty, and we let $p' := \langle 0, l' - 1 \rangle$.
 - Or wa is ambiguous and $(l' - 1)$ -faulty, so we let $p' := \langle ?, l' - 1 \rangle$. Notice that this can only happen if $l' \geq 2$.
 - Or wa is surely faulty, then $p' := \perp$ (this must happen if $l' = 1$).
- If $p = \langle ?, m \rangle$ with $m \leq l$, then w is ambiguous and m -faulty. The same three cases as before can arise:
 - Either wa is surely correct (only possible if $m \geq 2$), then it is still $(m - 1)$ -faulty, and we set $p' := \langle 0, m - 1 \rangle$.
 - Or wa is ambiguous and $(m - 1)$ -faulty, so we can set $p' := \langle ?, m - 1 \rangle$; this too requires $m \geq 2$.
 - Or wa is surely faulty, then $p' := \perp$ (this necessarily happens if $m = 1$).
- If $p = \perp$, then w is surely faulty, and so is wa , thus $p' := \perp$.

Observe that when the predictor moves from \top to some $\langle k, l' \rangle$, then (i) either $l' \leq l$ is finite, and both k' and l' can only decrease from there, so the predictions remain in $\mathbb{P}_{k,l}$, (ii) or $l' = l = \omega$, and then the predictions remain in $\mathbb{P}_{k,\omega}$ since $\omega - 1 = \omega$. Figure 3 abstractly illustrates the different types of compatible predictions used in the proof. \square




Fig. 3: Prediction transitions used in the proof of Proposition 17.

We remark that the proof of Proposition 17 does not try to optimize the prediction window. For instance, when $p = \langle k', l' \rangle$ with $k' > 0$, it may be possible that wa is in fact k'' - l'' -faulty for some $k'' \geq k' - 1$ and $k'' + l'' \leq k' - 1 + l'$. In this case Definition 13(iv) also allows to assign $\langle k'', l'' \rangle$ to p' instead. Since k'' may be larger than k' , the resulting prediction value may then be outside of $\mathbb{P}_{k,l}$.

Finally, we introduce the notion of *pilot* as an automata-based representation of k - l -active predictors. In Section 3 we will show how to find a finite-state pilot when \mathcal{A} is actively predictable and finite-state.

Definition 18 (pilot). *Let \mathcal{A} be an LTS, then $\mathcal{C} = \langle \mathcal{B}_{\mathcal{C}}, \text{cont}_{\mathcal{C}}, \text{pred}_{\mathcal{C}} \rangle$ is called pilot for \mathcal{A} over $\mathbb{P}' \subseteq \mathbb{P}$ if $\mathcal{B}_{\mathcal{C}} = \langle Q^c, q_0^c, \Sigma_o, T^c \rangle$ is a deterministic LTS with labellings $\langle \text{cont}_{\mathcal{C}}, \text{pred}_{\mathcal{C}} \rangle : Q^c \rightarrow 2^{\Sigma} \times \mathbb{P}'$. Let $h_{\mathcal{C}} = \langle \text{cont}, \text{pred} \rangle$ associated with \mathcal{C} be defined by $\text{cont}(w) = \text{cont}_{\mathcal{C}}(q)$ and $\text{pred}(w) = \text{pred}_{\mathcal{C}}(q)$ for all $w \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}))$, where q is the unique state such that $q_0^c \xrightarrow{w} q$. Then \mathcal{C} is a k - l -active predictor for \mathcal{A} if $h_{\mathcal{C}}$ is one.*

3 Controller construction

We solve the decision and synthesis problems simultaneously. We try to construct a pilot-based k - l -active predictor over some $\mathbb{P}' \subseteq \mathbb{P}$ for an LTS \mathcal{A} . The construction succeeds if and only if \mathcal{A} is k - l -actively predictable. According to Definition 13, the main challenges in building an active predictor are to ensure that (i) the controlled system remains live, (ii) the fault can be predicted at least k observations before its occurrences, and (iii) the prediction information is provided.

Our solution consists in building a turn-based game (see [12] for turn-based games) by taking into account the control that has already been performed.

Definition 19 (turn-based game). *A game \mathcal{G} with two players called Control and Environment is a tuple $\langle V_{\mathcal{C}}, V_{\mathcal{E}}, E, v_0, \text{WIN} \rangle$, where:*

- $V_{\mathcal{C}}, V_{\mathcal{E}}$ are the vertices owned by Control and Environment, respectively, and $V_{\mathcal{G}} = V_{\mathcal{C}} \uplus V_{\mathcal{E}}$ denoting all vertices, with $v_0 \in V_{\mathcal{C}}$ being an initial vertex;
- $E \subseteq V_{\mathcal{G}} \times V_{\mathcal{G}}$ is a set of directed edges such that for all $v \in V_{\mathcal{G}}$, there exists $(v, v') \in E$;
- $\text{WIN} \subseteq V_{\mathcal{G}}^{\omega}$ is a set of winning sequences.

Given a sequence $\rho = v_0 v_1 \dots v_n$, we denote $\rho[i] = v_i$. A *play* is a sequence of $V_{\mathcal{G}}^{\omega}$ such that $\rho[0] = v_0$ and $\langle \rho[i], \rho[i+1] \rangle \in E$ for all $i \geq 0$; we call $\rho^k := \rho[0] \dots \rho[k]$, for some $k \geq 0$, a *partial play* if $\rho[k] \in V_{\mathcal{C}}$, and define $\text{last}(\rho^k) := \rho[k]$. We write $\text{Play}^*(\mathcal{G})$ for the set of partial plays of \mathcal{G} . A play ρ is called *winning* (for Control) if $\rho \in \text{WIN}$.

A *Büchi game* $\langle V_{\mathcal{C}}, V_{\mathcal{E}}, E, v_0, V_{\mathcal{F}} \rangle$ defines a game $\langle V_{\mathcal{C}}, V_{\mathcal{E}}, E, v_0, \text{WIN} \rangle$ such that $\text{WIN} = \{ \rho \in V_{\mathcal{G}}^{\omega} \mid \rho[i] \in V_{\mathcal{F}} \text{ for infinitely many } i \}$. A *reachability game* $\langle V_{\mathcal{C}}, V_{\mathcal{E}}, E, v_0, V_{\mathcal{F}} \rangle$ defines a game $\langle V_{\mathcal{C}}, V_{\mathcal{E}}, E, v_0, \text{WIN} \rangle$ such that $\text{WIN} = V_{\mathcal{G}}^* V_{\mathcal{F}} V_{\mathcal{G}}^{\omega}$. A *safety game* $\langle V_{\mathcal{C}}, V_{\mathcal{E}}, E, v_0, V_{\mathcal{F}} \rangle$ defines a game $\langle V_{\mathcal{C}}, V_{\mathcal{E}}, E, v_0, \text{WIN} \rangle$ such that $\text{WIN} = V_{\mathcal{F}}^{\omega}$.

Definition 20 (strategy). *Let $\mathcal{G} = \langle V_{\mathcal{C}}, V_{\mathcal{E}}, E, v_0, \text{WIN} \rangle$ be a game. A strategy (for Control) is a function $\theta : \text{Play}^*(\mathcal{G}) \rightarrow V_{\mathcal{G}}$ such that $\langle \text{last}(\xi), \theta(\xi) \rangle \in E$ for all $\xi \in \text{Play}^*(\mathcal{G})$. A play ρ adheres to θ if $\rho[i] \in V_{\mathcal{C}}$ implies $\rho[i+1] = \theta(\rho^i)$ for all $i \geq 0$. A strategy is called winning if every play ρ that adheres to θ is winning. A positional (also called memoryless) strategy is a function $\theta' : V_{\mathcal{C}} \rightarrow V_{\mathcal{G}}$ such that $\langle v, \theta'(v) \rangle \in E$ for all $v \in V_{\mathcal{C}}$; we call θ' winning if the strategy θ with $\theta(\xi) = \theta'(\text{last}(\xi))$ is winning.*

To verify k - l -active predictability of a given system, the controller that we propose needs to memorize two subsets of states with the corresponding prediction information $\langle Q_c, Q_f, p \rangle$. The subset Q_c (resp. Q_f) represents the possible states reached by a correct (resp. faulty) run after the last observable action, and $Q_c \cup Q_f \neq \emptyset$. The prediction information $p \in \mathbb{P}'$ is (non-deterministically) decided based on the current observations. We denote $\text{Reach}(\langle Q_c, Q_f, p \rangle) := Q_c \cup Q_f$ and $\tilde{Q} := 2^Q \setminus \{\emptyset\}$. The set of possible tuples memorized by the controller is defined as $S_{\mathbb{P}'} = S_{\mathbb{P}'}^c \cup S_{\mathbb{P}'}^a \cup S_{\mathbb{P}'}^f$, where:

- $S_{\mathbb{P}'}^c = \tilde{Q} \times \{\emptyset\} \times \{ p \in \mathbb{P}' \mid \kappa(p) \geq 0 \}$

- $S_{\mathbb{P}'}^a = \tilde{Q} \times \tilde{Q} \times (\mathbb{P}' \cap (\{?\} \times \mathbb{N}_\omega^+))$
- $S_{\mathbb{P}'}^f = \{\emptyset\} \times \tilde{Q} \times \{\perp\}$

In the following, we will simply write S for $S_{\mathbb{P}'}$ when \mathbb{P}' is clear from context.

The controller needs to update the state subsets after an observable action, for which we first define some sets of possible next states from a given state q after $a \in \Sigma_o$.

- $NO_{\mathcal{A}}(q, a) = \{q' \mid q \xrightarrow{\sigma}_{\mathcal{A}} q', \sigma \in \Sigma_{uo}^* a\}$
- $NOC_{\mathcal{A}}(q, a) = \{q' \mid q \xrightarrow{\sigma}_{\mathcal{A}} q', \sigma \in (\Sigma_{uo} \setminus \{f\})^* a\}$
- $NOF_{\mathcal{A}}(q, a) = \{q' \mid q \xrightarrow{\sigma}_{\mathcal{A}} q', \sigma \in \Sigma_{uo}^* f \Sigma_{uo}^* a\}$

One can omit the subscript \mathcal{A} when there is no ambiguity. The extension to a set of states is defined in a natural way, e.g. $NO(Q', a) = \bigcup_{q \in Q'} NO(q, a)$. We now define how the controller updates its tuple once an observable action occurs. In the following, \emptyset represents a state in which the controller has lost, and we denote $S^\emptyset := S \cup \{\emptyset\}$.

Definition 21 (knowledge update). *Let \mathcal{A} be an LTS, $\mathbb{P}' \subseteq \mathbb{P}$, and $k \geq 0$. Then the knowledge transition relation $\Delta_{\mathcal{A}}^k \subseteq S \times \Sigma_o \times S^\emptyset$ is defined as follows. Let $s = \langle Q_c, Q_f, p \rangle \in S$ and $a \in \Sigma_o$. Then $\langle s, a, s' \rangle \in \Delta_{\mathcal{A}}^k$ if and only if:*

1. either $s' = \langle NOC(Q_c, a), NOF(Q_c, a) \cup NO(Q_f, a), p' \rangle \in S$ and $\langle p, p' \rangle$ are k - l -compatible;
2. or $s' = \emptyset$ when there is no $s'' \in S$ such that $\langle s, a, s'' \rangle \in \Delta_{\mathcal{A}}^k$.

Notice that, given s and a , the choice of s' is largely deterministic except for p' , which must be k - l -compatible with p . When s' has no prediction consistent with the updated correct resp. faulty state subsets, cf Definition 13(iii), then the only possible update is to \emptyset .

Example 22. Consider the LTS in Figure 1 and assume that $\Sigma_1 = \{a, c\}$, $\Sigma_2 = \{a, b\}$ and $\Sigma_c = \{a, b, c\}$.

1. Let $s = \langle \{q_0\}, \emptyset, \top \rangle$. If the observable action a is chosen, then we have $\langle s, a, s' \rangle \in \Delta_{\mathcal{A}}^k$, where $s' = \langle \{q_1, q_4\}, \emptyset, \top \rangle$. Notice that $\langle \top, \top \rangle$ are k - l -compatible.
2. Let $s = \langle \{q_2, q_5\}, \emptyset, \top \rangle$ after observing a and d . If a is chosen from here, we can only have $\langle s, a, \emptyset \rangle \in \Delta_{\mathcal{A}}^k$. The reason is that after a , the system can end up in either q_3 (with a fault) or in q_5 (without fault), the next prediction should thus be an ambiguous one, i.e., $\langle ?, m \rangle$. However, $\langle \top, \langle ?, m \rangle \rangle$ are not k - l -compatible. It follows that there does not exist $s'' \in S$ such that $\langle s, a, s'' \rangle \in \Delta_{\mathcal{A}}^k$. Hence we have $\langle s, a, \emptyset \rangle \in \Delta_{\mathcal{A}}^k$ by Definition 21.

The objective of Control is to obtain a winning play by suitably restricting the possible actions, and any winning strategy corresponds to a controller with which the controlled system is predictable. The game begins with Control to choose a prediction for ε . Then the game proceeds in rounds: 1) Control restricts the set of possible actions to some Σ' ; 2) Environment chooses $a \in \Sigma'$ to determine the next state. 3) Control updates its knowledge.

The choices of Control are subject to some restrictions. Indeed, each state $s = \langle Q_c, Q_f, p \rangle$ represents Control's knowledge about the current potential states of \mathcal{A} as well as the corresponding prediction information. To ensure that the controlled system remains live, the set of possible actions Σ' must not cause deadlocks in any state reachable by unobservable actions from $Q_c \cup Q_f$. Also, Control cannot prevent the uncontrollable actions. So we define the admissible sets and the game as follows, where we use $\Sigma_{PO}(q) = \{a \in \Sigma_o \mid q \xrightarrow{\sigma} q'', \sigma \in \Sigma_{uo}^* a\}$ to denote the possible next observable actions from the state q , which can be extended to a set of states in a natural way.

Definition 23 (admissible action set). *Let $\mathcal{A} = \langle Q, q_0, \Sigma, T \rangle$ be an LTS and $Q' \subseteq Q$ be a subset of states. We call $\Sigma' \subseteq \Sigma_o$ an admissible set for Q' if it fulfills the following conditions:*

- $\Sigma_{uco} \subseteq \Sigma'$ as any action in Σ_{uco} is observable but not controllable.
- for all $q' \in Q'$, $q \in Q$, and $\sigma \in \Sigma_{uo}^*$, $q' \xrightarrow{\sigma} q$ implies $\Sigma_{PO}(q) \cap \Sigma' \neq \emptyset$.

Proposition 27. *Given $h = \langle cont, pred \rangle$ a k - l -active predictor over \mathbb{P}' for an LTS \mathcal{A} , there exists a corresponding winning strategy θ_h in the game $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$.*

Proof.

- We define a strategy θ_h based on h and demonstrate its existence in $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$. We also show that any play ξ adhering to θ_h and all its prefixes possess the following five invariants, where $Obs(\xi)$ denotes the observable actions along a partial play ξ . Formally, $Obs(\varepsilon) = \varepsilon$, $Obs(\xi \langle s, a \rangle) = Obs(\xi)a$ for $\langle s, a \rangle \in V_2$, and $Obs(\xi r) = Obs(\xi)$ for $r \notin V_2$.

- $\phi_{-\circ}$: ξ never enters \circ ;
- ϕ_o : $Obs(\xi) \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$;
- ϕ_c : if $last(\xi) = \langle Q_c, Q_f, p \rangle \in S$, then $Q_c = NOC_{\mathcal{A}_{cont}}(q_0, Obs(\xi))$;
- ϕ_f : if $last(\xi) = \langle Q_c, Q_f, p \rangle \in S$, then $Q_f = NOF_{\mathcal{A}_{cont}}(q_0, Obs(\xi))$;
- ϕ_p : if $last(\xi) = \langle Q_c, Q_f, p \rangle \in S$, then $p = pred(Obs(\xi))$.

It is trivial to show that these five invariants are initially true when $\xi = \iota$. From Definition 25, a play can enter \circ only by a transition of E_3 . It is thus enough to prove $\phi_{-\circ}$ on all transitions in ξ belonging to E_3 . We show that ξ preserves ϕ_o on transitions of E_2 , where $Obs(\xi)$ changes. Similarly, ϕ_c , ϕ_f and ϕ_p should be kept true on the transitions of E_l and E_3 , where their antecedent is true. Now we define θ_h by considering three subsets of Control vertices, for each one their existence in $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$ and invariant preservation being demonstrated as well.

1. $\theta_h(\iota) = \langle \{q_0\}, \emptyset, pred(\varepsilon) \rangle$;
 - Since h is a k - l -active predictor, $pred(\varepsilon)$ is thus k - l -initial by Definition 13 ((ii)). Hence, $\langle \iota, \theta_h(\iota) \rangle \in E_l$ from Definition 25.
 - For $\xi = \iota \theta_h(\iota)$, ϕ_c , ϕ_f and ϕ_p are true as $Obs(\xi) = \varepsilon$.
 2. Given any $\xi \in Play^*(\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l})$ adhering to θ_h with $last(\xi) = s \in S$, for which all invariants are true, we have $\theta_h(\xi) = \langle s, \Sigma' \rangle$, where $\Sigma' = cont(Obs(\xi)) \setminus \Sigma_{uo}$;
 - $\Sigma_{uco} \cup \Sigma_{uo} \subseteq cont(w)$ for any $w \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}))$ due to Definition 9, thus $\Sigma_{uco} \subseteq \Sigma'$. \mathcal{A}_{cont} is live by Definition 13 ((i)), i.e., no deadlock in the controlled system, thus $\Sigma' \in adm(Reach(s))$ by Definition 23. Hence, $\langle s, \theta_h(\xi) \rangle \in E_1$.
 - We show that ϕ_o keeps true for any play $\xi' = \xi \theta_h(\xi) \langle s, a \rangle$, where $\langle \theta_h(\xi), \langle s, a \rangle \rangle \in E_2$. Since ϕ_o is true for ξ , let $\sigma \in \mathcal{L}^*(\mathcal{A}_{cont})$, $\sigma \in \Sigma^* \Sigma_o$, such that $\mathcal{P}(\sigma) = Obs(\xi)$. For any $\langle \theta_h(\xi), \langle s, a \rangle \rangle \in E_2$, we have $a \in \Sigma_{PO}(Reach(s)) \cap \Sigma'$ and $\Sigma' \subseteq cont(Obs(\xi))$. It follows that there exists $\sigma' \in \Sigma_{uo}^*$ such that $\sigma \sigma' a \in \mathcal{L}^*(\mathcal{A}_{cont})$, therefore $Obs(\xi') = Obs(\xi)a = \mathcal{P}(\sigma)a = \mathcal{P}(\sigma \sigma' a) \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$. Hence ϕ_o is true for ξ' .
 3. Given any $\xi \in Play^*(\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l})$ adhering to θ_h with $last(\xi) = \langle \langle Q_c, Q_f, p \rangle, a \rangle \in V_2$, for which all invariants are true, we have $\theta_h(\xi) = \langle Q'_c, Q'_f, pred(Obs(\xi)) \rangle$, where $Q'_c = NOC(Q_c, a)$ and $Q'_f = NOF(Q_c, a) \cup NO(Q_f, a)$.
 - We first show that $\langle p, pred(Obs(\xi)) \rangle$ is k - l -compatible. Let $\xi' \in Play^*(\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l})$ and $last(\xi') = \langle Q_c, Q_f, p \rangle$. We denote $Obs(\xi') = w$, and thus we have $Obs(\xi) = wa$. From the assumption that ϕ_o and ϕ_p are true for ξ and its prefixes, we have $w, wa \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$, $p = pred(w)$, and $pred(Obs(\xi)) = pred(wa)$. Since h is a k - l -active predictor, $\langle p, pred(Obs(\xi)) \rangle$ is k - l -compatible by Definition 13 ((iv)). Hence, we can infer from the construction of Q'_c and Q'_f that $\langle last(\xi), \theta_h(\xi) \rangle \in E_3$.
 - We need to show all invariants, except ϕ_o , are true for $\xi'' = \xi \theta_h(\xi)$. First, since $\phi_{-\circ}$ is true for ξ and $\langle last(\xi), \theta_h(\xi) \rangle \in E_3$ shown as above, where $\theta_h(\xi) \in S$, then $\phi_{-\circ}$ is also true for ξ'' by Definition 21. Now consider ϕ_c . We have $Q_c = NOC_{\mathcal{A}_{cont}}(q_0, Obs(\xi'))$ as ϕ_c is true for the prefixes of ξ . From $Q'_c = NOC(Q_c, a)$ and $wa = Obs(\xi')a \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$, we have $Q'_c = NOC_{\mathcal{A}_{cont}}(q_0, Obs(\xi')a)$, thus $Q'_c = NOC_{\mathcal{A}_{cont}}(q_0, Obs(\xi''))$. So ϕ_c is true for ξ'' . ϕ_f can be proved to be true for ξ'' in a similar way. ϕ_p is true for ξ'' as $Obs(\xi) = Obs(\xi'')$.
- We next show that θ_h is winning. Given a play $\rho \in V_{\mathcal{G}}^\omega$ adhering to θ_h , we have $\rho[1] = \langle \{q_0, \emptyset, p_0\} \rangle$, p_0 is k - l -initial. We prove that ρ is winning when $p_0 = \top$ or $p_0 = \langle k', l' \rangle$.

- If $p_0 = \top$, there are two cases as ρ never enters \circlearrowleft from $\phi_{-\circlearrowleft}$. One is that for all $s = \langle Q_c, Q_f, p \rangle \in \rho$ and $s \in S$, $p = \top$. In other words, we have infinitely $s \in \tilde{Q} \times \{\emptyset\} \times \{\top\} \subseteq V_F$ in ρ , which is thus winning. Consider the second case, where there exists $s = \langle Q_c, Q_f, p \rangle \in \rho$ and $s \in S$ such that $p \neq \top$. Let $\rho[i] = \langle \langle Q_c, \emptyset, \top \rangle, a \rangle$, and $\rho[i+1] = \langle Q_c, Q_f, p \rangle$, where $p \neq \top$. From $\langle \rho[i], \rho[i+1] \rangle \in E_3$, we can infer that $\langle \top, p \rangle$ is k - l -compatible, thus $\kappa(p) \geq k$, so $p = \langle k', l' \rangle \in \mathbb{N} \times \mathbb{N}_\omega^+$. We analyze two cases based on l' .
 1. $l' < \omega$: from the invariants ϕ_o and ϕ_p , we have $Obs(\rho^{i+1}) \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$ and $pred(Obs(\rho^{i+1})) = p$. It follows that $Obs(\rho^{i+1})$ is k' - l' -faulty, thus $k'+l'$ faulty in \mathcal{A}_{cont} as h is a k - l -active predictor. One can thus infer with the same invariants that there exists $m < \omega$ such that $\rho[m] \in \{\emptyset\} \times \tilde{Q} \times \{\perp\}$. From ϕ_c , one can deduce that for any $\rho[m'] \in S$, $m' > m$, $\rho[m'] \in \{\emptyset\} \times \tilde{Q} \times \{\perp\} \subseteq V_F$, thus ρ is winning.
 2. $l' = \omega$: as above, one can show that $Obs(\rho^{i+1})$ is ω -faulty, thus m -faulty for some $m \in \mathbb{N}$ by Definition 4. Hence, ρ is winning with the same reasoning.
- If $p_0 = \langle k', l' \rangle$, one can demonstrate that ρ is winning exactly in the same way for the second case when $p_0 = \top$.

□

The existence of a winning strategy implies the existence of a positional one due to well-known results of game theory (see e.g. [12] for all results here related to turn-based games). For the reverse direction, we next define a pilot from a positional winning strategy in $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k,l}$ before proving that this pilot is a k - l -active predictor.

Definition 28. Let θ be a positional winning strategy in $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k,l}$. We define a pilot $\mathcal{C}_\theta := \langle \mathcal{B}_\theta, cont_\theta, pred_\theta \rangle$ over \mathbb{P}' as follows:

- $\mathcal{B}_\theta = \langle Q^\theta, q_0^\theta, \Sigma_o, T^\theta \rangle$, where
 1. $Q^\theta = \{q \in S \mid q = last(\xi_\theta) \text{ and } \xi_\theta \in Play^*(\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k,l}) \text{ adhering to } \theta\}$
 2. $q_0^\theta = \theta(\iota)$
 3. $T^\theta(s, a) = \theta(\langle s, a \rangle)$
- $cont_\theta(s) = \Sigma' \cup \Sigma_{uo}$ for any $s \in Q^\theta$, where $\theta(s) = \langle s, \Sigma' \rangle$;
- $pred_\theta(s) = p$, for any $s = \langle Q_c, Q_f, p \rangle \in Q^\theta$

Proposition 29. Let θ be a positional winning strategy in $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k,l}$. Then \mathcal{C}_θ is a k - l -active predictor over \mathbb{P}' for \mathcal{A} .

Proof. For any $s \in Q^\theta$, $cont_\theta(s) = \Sigma' \cup \Sigma_{uo}$, which contains thus Σ_{uo} . From Definition 23, $cont_\theta(s)$ contains also Σ_{ucco} as $\theta(s) = \langle s, \Sigma' \rangle$. Hence, due to Definition 18, the corresponding $h_{\mathcal{C}_\theta} = \langle cont, pred \rangle$ is constructed from \mathcal{C}_θ , where $cont$ is a controller by Definition 9.

Now we show that $h_{\mathcal{C}_\theta}$ satisfies all conditions of Definition 13, which means that it is a k - l -active predictor, and thus is \mathcal{C}_θ .

- (i) We show that given any $\sigma \in \mathcal{L}^*(\mathcal{A}_{cont})$ and $w = \mathcal{P}(\sigma)$, $\exists s \in Q^\theta$ such that $q_0^\theta \xrightarrow{w}_{\mathcal{B}_\theta} s$, which is demonstrated in an inductive way on w (see the following), and thus we have $cont(w) = cont_\theta(s)$. Now we have $\theta(s) = \langle s, cont(w) \setminus \Sigma_{uo} \rangle$, therefore $cont(w) \setminus \Sigma_{uo} \in adm(Reach(s))$ from E_1 . It follows that there exists $a \in cont(w)$ such that $\sigma a \in \mathcal{L}^*(\mathcal{A}_{cont})$. Hence, \mathcal{A}_{cont} is live.
 - Let $w = \varepsilon$, then $q_0^\theta \xrightarrow{w}_{\mathcal{B}_\theta} q_0^\theta$, which is the base case.
 - For the induction step, suppose $w \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$ and $q_0^\theta \xrightarrow{w}_{\mathcal{B}_\theta} s$. For any $wa \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$, we show that $\exists s'$ such that $q_0^\theta \xrightarrow{wa}_{\mathcal{B}_\theta} s'$. From the assumption, we have $\theta(s) = \langle s, \Sigma' \rangle$ such that $a \in \Sigma'$. $wa \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$ implies also $a \in \Sigma_{PO}(Reach(s))$. Thus, $\langle \langle s, \Sigma' \rangle, \langle s, a \rangle \rangle \in E_2$ and $\exists s' \in S$ such that $\theta(\langle s, a \rangle) = s'$. Hence, we have $T^\theta(s, a) = \theta(\langle s, a \rangle) = s'$, thus $q_0^\theta \xrightarrow{wa}_{\mathcal{B}_\theta} s'$.
- (ii) $pred(\varepsilon) = pred_\theta(q_0^\theta)$ is k - l -initial, deduced directly from E_1 .
- (iii) Let $w \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{cont}))$, and $q_0^\theta \xrightarrow{w}_{\mathcal{B}_\theta} s$. One thus has $cont(w) = cont_\theta(s)$ and $pred(w) = pred_\theta(s)$.

- If $\text{pred}(w) = \top$, we show that w is $k + 1$ -correct in $\mathcal{A}_{\text{cont}}$. Let $wa \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{\text{cont}}))$. From above, there exists $s' \in S$ such that $q_0^\theta \xrightarrow{wa} \mathcal{B}_\theta s'$, therefore $\text{pred}(wa) = \text{pred}_\theta(s')$. Furthermore, since $\kappa(\text{pred}(w)) = \omega + 1$, so $\kappa(\text{pred}(wa)) \geq k$, which implies $\text{pred}(wa) = \top$ or $\text{pred}(wa) = \langle k', l' \rangle$, where $k' \geq k$. It follows that wa is k' -correct in $\mathcal{A}_{\text{cont}}$ (see the next step), thus also k -correct. Hence, w is $k + 1$ -correct.
- If $\text{pred}(w) = \langle k', l' \rangle$, we prove that w is k' -correct and $k' + l'$ -faulty in $\mathcal{A}_{\text{cont}}$ by induction on k' and $k' + l'$, respectively.

The base case for k' -correct is 0-correct. For any $\text{pred}(w) = \text{pred}_\theta(s) = \langle k', l' \rangle$, $k' \geq 0$, we have that w is surely correct since $s \in S_{\mathbb{P}'}^c$, thus 0-correct in $\mathcal{A}_{\text{cont}}$ from Definition 4. For the induction step, let $k' \in \mathbb{N}$. Suppose that for any w such that $\text{pred}(w) = \langle k'', l'' \rangle$ and $k'' \geq k'$, w is k' -correct in $\mathcal{A}_{\text{cont}}$. We show that w' is $k' + 1$ -correct in $\mathcal{A}_{\text{cont}}$ when $\text{pred}(w') = \langle k' + 1, l_1 \rangle$, whose κ is $k' + 1$. From k - l -compatibility, we have $\kappa(\text{pred}(w'a)) \geq k'$ for any $w'a \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{\text{cont}}))$. From the assumption, $w'a$ is k' -correct in $\mathcal{A}_{\text{cont}}$. Thus, w' is $k' + 1$ -correct in $\mathcal{A}_{\text{cont}}$.

For $k' + l'$ -faulty, we prove it in two cases.

- (a) When $l' < \omega$, the base case is 1-faulty. Let $\text{pred}(w) = p = \langle 0, 1 \rangle$, thus $\mu(p) = 1$. For any $wa \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{\text{cont}}))$, we have $\mu(\text{pred}(wa)) = 0$ and $\text{pred}(wa) = \perp$, which means that wa is surely faulty in $\mathcal{A}_{\text{cont}}$ (see last step), and thus w is 1-faulty in $\mathcal{A}_{\text{cont}}$. Given any w' such that $\mu(\text{pred}(w')) \leq 1$, w' is 1-faulty. For the induction step, let $n \in \mathbb{N}$. Suppose that for any w such that $\text{pred}(w) = \langle k', l' \rangle$, where $k' + l' \leq n$ ($\mu(\text{pred}(w)) \leq n$), w is n -faulty in $\mathcal{A}_{\text{cont}}$. We next show that w' is $n + 1$ -faulty in $\mathcal{A}_{\text{cont}}$ when $\text{pred}(w') = \langle k'', l'' \rangle$, where $k'' + l'' = n + 1$. Again from k - l -compatibility, we have $\mu(\text{pred}(w'a)) \leq n$ for any $w'a \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{\text{cont}}))$. From the assumption, $w'a$ is n -faulty in $\mathcal{A}_{\text{cont}}$. Thus, w' is $n + 1$ -faulty in $\mathcal{A}_{\text{cont}}$.
 - (b) When $l' = \omega$, we prove that w is ω -faulty, i.e., $\exists m \in \mathbb{N}$ such that w is m -faulty. To this end, we define $v_w \in S$, the vertex reached in $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$ after observing w in a play conforming to the strategy θ . Consider \mathcal{G}' , the subgraph of $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$ consisting of successors of v_w when following θ without the vertices whose prediction is \perp . \mathcal{G}' is acyclic since otherwise there would an infinite play conforming to the strategy θ where predictions are neither \top nor \perp , thus losing. Thus after a finite number of steps (bounded by $3|S|$), a path conforming to θ and starting from v_w must exit \mathcal{G}' . Since it can neither reach \circlearrowleft or a state with prediction \top , it must reach a state in $\{\emptyset\} \times \tilde{Q} \times \{\perp\}$ which corresponds to a faulty run in $\mathcal{A}_{\text{cont}}$. Thus w is m -faulty for some $m \leq |S| + 1$.
 - If $\text{pred}(w) = \langle ?, m \rangle$, one can demonstrate that w is m -faulty in the same way as above for $k' + l'$ -faulty. Furthermore, $s \in S_{\mathbb{P}'}^g$, implying the ambiguity of w .
 - If $\text{pred}(w) = \perp$, then $s \in S_{\mathbb{P}'}^f$. Thus, w is surely faulty by the definition.
- (iv) Let $a \in \Sigma_o$, $w, wa \in \mathcal{P}(\mathcal{L}^*(\mathcal{A}_{\text{cont}}))$. From Definition 21, $\langle \text{pred}(w), \text{pred}(wa) \rangle$ is k - l -compatible since the corresponding strategy is winning, not entering \circlearrowleft .

□

Combining the results of Propositions 27 and 29, we obtain that the active-predictability problem for an LTS \mathcal{A} with n states reduces to solving a Büchi game with $2^{\mathcal{O}(n)}$ vertices. Since Büchi games can be solved in polynomial time, we obtain the following result:

Theorem 30. *The active-predictability problem for finite-state LTS belongs to EXPTIME.*

We conclude the section with a supplementary result showing that due to the special structure of $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$, it can actually be solved in linear time (w.r.t. the size of the game), and not in quadratic time as performed for general Büchi games.

Proposition 31. *If \mathcal{A} is a finite-state LTS and $\mathbb{P}' \subseteq \mathbb{P}$, then $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$ can be solved in $\mathcal{O}(|E|)$.*

Proof. It is well-known that the solution of a Büchi game generally takes quadratic time. However, the particular Büchi games we generate for the fault prediction problem can be reduced to solving reachability game and then a safety game deduce from the solution of the reachability game. Since these types of games can be solved in linear time, so can $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$.

The winning vertices of the Büchi game $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l} = \langle V_C, V_E, E, \iota, V_F \rangle$ can be partitioned into $V_F = V_F^S \cup V_F^R$, with $V_F^S := \tilde{Q} \times \{\emptyset\} \times \{\top\}$ and $V_F^R := \{\emptyset\} \times \tilde{Q} \times \{\perp\}$ (cf Definition 25). Thus, a winning play either never admits a fault and predicts \top all the time, or it eventually passes to a different prediction, whence it must eventually reach the prediction \perp , having definitely committed a fault. Once the prediction \perp is reached, it will remain forever (cf Figure 3), so reachability is equivalent to infinitely repeated reachability for these vertices. A customized algorithm for $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$ is therefore as follows:

1. First solve the reachability game $\mathcal{G}_R := \langle V_C, V_E, E, \iota, V_F^R \rangle$ and determine the winning states of \mathcal{G}_R belonging to S , say S' . Let us denote θ_1 the (positional) strategy corresponding to the winning states.
2. Then transform $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$ as follows. All vertices different from ι , not associated with prediction \top and not belonging to S' are deleted and their incoming edges are redirected to \emptyset . The vertices that belong to S' are made absorbing. Then in this arena, solve the safety game that consists in avoiding \emptyset . Let us denote θ_2 the (positional) strategy corresponding to the winning states.

Then the controller has a winning strategy in $\mathcal{G}_{\mathcal{A}, \mathbb{P}'}^{k, l}$ if ι belongs to the winning states of the safety game. The corresponding (positional) strategy consists in playing according to θ_2 as long as the play do not enter S' and when possibly entering in S' consists in playing according to θ_1 . □

4 Bound analysis

We first prove that it is EXPTIME-hard to decide whether a given LTS \mathcal{A} is actively k - l -predictable, independently of k and l . The proof is similar to the proof in [13] that active *diagnosability* is EXPTIME-hard and relies on a reduction from safety games with imperfect information [3].

Theorem 32. *The active-predictability decision problem is EXPTIME-hard.*

Proof. A tuple $\mathcal{G} = \langle L, l_0, \Sigma, \Delta, O, F, obs \rangle$ is called a *safety game with imperfect information*, where L is a finite set of locations with initial location $l_0 \in L$, Σ is a finite alphabet, $\Delta \subseteq L \times \Sigma \times L$ is the transition relation such that for all $l \in L$ and $a \in \Sigma$ there exists at least one l' with $\langle l, a, l' \rangle \in \Delta$, O is a finite set of observations, $F \subseteq O$ are the final observations, and $obs : L \mapsto O$ is the observation mapping.

\mathcal{G} is a turn-based game played by two players called Control and Environment. It starts in location l_0 with Control to play. In the first round, Control chooses a letter a_0 in Σ , and then Environment chooses a location l_1 such that $\langle l_0, a_0, l_1 \rangle \in \Delta$. Control only observes $o_1 = obs(l_1)$. The next rounds are played similarly. Control wins if for all i , $o_i \notin F$.

Figure 5 (a) shows an example of a game with alphabet $\Sigma = \{a, b\}$ and observations $O = \{\mathbf{o}, \mathbf{p}, \mathbf{q}\}$, annotated next to the locations, where $\mathbf{q} \in F$ is the only final observation. Control must therefore prevent the system from entering location l_3 .

Verifying the existence of a winning strategy for Control is EXPTIME-complete [3]. We now describe the reduction of this problem to an active-predictability decision problem with LTS \mathcal{A} defined as follows.

- Q , the set of states, is defined by $Q = L \uplus ((L \setminus obs^{-1}(F)) \times \Sigma) \uplus \{\perp\}$ and $q_0 = l_0$.
- The alphabet is $\Sigma' = \Sigma \uplus O \uplus \{u, f, z\}$. The unobservable events are u and f and the (observable) uncontrollable events are $O \uplus \{z\}$.
- T , the transition relation, is defined as follows.
 1. For all $l \in L \setminus obs^{-1}(F)$ and $a \in \Sigma$, $\langle l, a, \langle l, a \rangle \rangle \in T$.
 2. For all $l \in L \setminus obs^{-1}(F)$, $a \in \Sigma$ and $l' \in L$, $\langle \langle l, a \rangle, obs(l'), l' \rangle \in T$ if $\langle l, a, l' \rangle \in \Delta$.
 3. $\langle \perp, z, \perp \rangle \in T$, and for all $l \in obs^{-1}(F)$, $\langle l, u, \perp \rangle$ and $\langle l, f, \perp \rangle$ belong to T .

Evidently, \mathcal{A} is actively predictable if one can construct a finite-state pilot that avoids $obs^{-1}(F)$; if that is possible, then no fault can ever occur, and the prediction is always \top . Otherwise, the controlled system has a choice of going to z with a fault in the next step, or going there with u and never committing a fault, so




Fig. 5: (a) A safety game with imperfect information; Control must avoid observation q ; (b) the corresponding active-predictability problem constructed by Theorem 32.

predicting a fault becomes impossible. In addition, such a pilot only “controls” the subset of states $L \setminus \text{obs}^{-1}(F)$ and due to the assumptions on \mathcal{G} , liveness remains ensured when restricting the allowed events to a singleton. Furthermore the information available to the pilot is exactly that of Control, i.e. the letters chosen by Control himself and the observations due to the states chosen by Environment. Therefore, a winning strategy for Control in \mathcal{G} provides an active predictor for \mathcal{A} and vice versa. Figure 5 (b) shows the LTS constructed from the safety game in Figure 5 (a). \square

Together with Theorem 30, we obtain the following corollary.

Corollary 33. *The active-predictability decision problem is EXPTIME-complete.*

We study the relation between active predictability and active safety. Theorem 34 relates the maximal advance warning for fault predictions to the number of states in \mathcal{A} .

Theorem 34. *Let \mathcal{A} be an LTS with n states. If \mathcal{A} is 2^n -active-predictable, then it is actively safe.*

Proof. If \mathcal{A} is 2^n - ω -active-predictable then by definition there exists a 2^n - ω -active predictor $h = \langle \text{cont}, \text{pred} \rangle$ over $\mathbb{P}' := \mathbb{P}_{k,\omega}$ for \mathcal{A} , and by Proposition 27 there exists a winning strategy θ in $\mathcal{G}_{\mathcal{A},\mathbb{P}'}^{k,\omega}$. In turn, this winning strategy provides a pilot $\mathcal{C}_\theta = \langle \mathcal{B}, \text{cont}', \text{pred}' \rangle$ according to Proposition 29; let $\mathcal{B} = \langle Q, q_0, \Sigma_o, T \rangle$. We shall construct a new pilot \mathcal{C} for \mathcal{A} over $\{\top\}$, proving that \mathcal{A} is actively safe.

Remember that Q is the set of Controller-owned vertices in $\mathcal{G}_{\mathcal{A},\mathbb{P}'}^{k,\omega}$ that can be reached by plays adhering to θ and that these vertices are a subset of $S_{\mathbb{P}'}$. For $q, q' \in Q$, let us write $q \prec q'$ if q' is reachable from q in \mathcal{B} . Since θ is positional and winning, \prec must be an acyclic relation between those states of Q that are not members of V_F , i.e. their associated prediction is neither \top nor \perp (cf Definition 25). We now call $q \in Q$ a *cutoff* if q is of the form $\langle Q_c, Q_f, p \rangle$ and there exists a state $q' = \langle Q_c, Q_f, p' \rangle$ with $p' \neq p$ and $q' \prec q$. Let $\text{co}(q)$, the *corresponding state* of q , denote the state that is \prec -minimal among all the choices for q' ; due to the structure of the states outside V_F , $\text{co}(q)$ is unique and not a cutoff itself. Moreover, a state of Q is called *useless* if it is either a cutoff or all its (immediate) predecessors in \mathcal{B} are useless, and *useful* otherwise.

Remember that $S_{\mathbb{P}'}$ is a union of $S_{\mathbb{P}'}^c$, $S_{\mathbb{P}'}^a$, and $S_{\mathbb{P}'}^f$, where $S_{\mathbb{P}'}^c$ contains the states of the form $\langle Q_c, \emptyset, p \rangle$, with $\kappa(p) \geq 0$. Thus, states in $S_{\mathbb{P}'}^c$ are only reached through correct runs in $\mathcal{A}_{\text{cont}'}$, due to invariants ϕ_c and ϕ_f in the proof of Proposition 27. Let $S' := \{ \langle Q_c, \emptyset, p \rangle \mid \kappa(p) = 0 \}$. It follows from the construction of $\mathcal{G}_{\mathcal{A},\mathbb{P}'}^{k,\omega}$ (cf Definition 21 and Definition 25) that any path from q_0 to a state from S' is of length at least 2^n , so by pigeonhole principle, any path leading to S' contains a cutoff. Since $S_{\mathbb{P}'}^a \cup S_{\mathbb{P}'}^f$ can only be reached by going through S' , those states are useless.

We can now construct the desired pilot \mathcal{C} by “folding” cutoffs back onto their corresponding states. We remark in this context that $\text{Reach}(q) = \text{Reach}(\text{co}(q))$, and therefore the admissible control choices for both

states are the same; proving that the resulting controlled system is live depends only on this property, cf the proof of Proposition 29 (i). Since the controlled system never admits a fault, the prediction can be \top in all cases. More formally, $\mathcal{C} := \langle \langle Q', q_0, \Sigma_o, T' \rangle, cont', pred'' \rangle$, where Q' is the useful subset of Q , and for all $q \in Q'$, $a \in \Sigma_o$:

- $T'(q, a) = T(q, a)$ if $T(q, a) \in Q'$ and $T'(q, a) = co(T(q, a))$ otherwise;
- $pred''(q) = \top$.

□

Theorem 34 implies that if a system is not actively safe, then there is an exponential upper bound on the advance warning that an active predictor can issue. This bound is asymptotically precise, as the following family of examples shows.

Theorem 35. *There exists a family of systems $(\mathcal{A}_n)_{n \geq 1}$ with $\mathcal{O}(n)$ states such that \mathcal{A}_n is not actively safe but 2^n -active-predictable.*




Fig. 6: A 2^n -active predictable LTS with $\mathcal{O}(n)$ states, where $\Sigma_o = \Sigma_c = \{0, \dots, n\}$, $\Sigma_i = \{i + 1, \dots, n\}$, and $\Gamma_i = \{0, \dots, i - 1\}$.

Proof. Figure 6 shows a family of LTS with $\mathcal{O}(n)$ states but an alphabet of size $\mathcal{O}(n)$ and $\mathcal{O}(n^2)$ transitions. We first provide a proof for this family as it is easier to understand. After this, we provide a more complex example with a constant-size alphabet and $\mathcal{O}(n)$ states and transitions.

Variable-size alphabet Consider the LTS shown in Figure 6. The observable actions are $\{0, \dots, n\}$, all of which are controllable. There are only two unobservable actions, u and the fault f . We abbreviate by $\Sigma_i := \{i + 1, \dots, n\}$ the actions larger than i for $0 \leq i < n$, and by $\Gamma_i := \{0, \dots, i - 1\}$ the actions smaller than i for $0 < i \leq n$.

The initial state is p . Evidently \mathcal{A}_n is actively safe if a controller can avoid both p' and q ; as we shall see, this is impossible. However, the system is actively predictable if the controller can at least avoid q . We shall see that this is indeed possible while entering p' only after 2^n steps, by simulating a binary counter.

We can assume (w.l.o.g.) that the controller permits a single action from Σ_o in each step and hence the controlled system will admit a single infinite observation sequence ρ . Having allowed a prefix σ of ρ , let $R(\sigma)$ be the set of states that this sequence can lead to. If the controller wants to keep the system from making a fault, it must ensure that $R(\sigma)$ remains within the set $R := \{p, s_0, \dots, s_{n-1}\}$. When $R(\sigma) \subseteq R$, let us associate a measure defined as $I(\sigma) := \sum_{s_i \in R(\sigma)} 2^i$. We observe the following:

- $R(\varepsilon) = \{p\}$, hence $I(\varepsilon) = 0$.
- If $s_i \in R(\sigma)$, then the controller must not allow action i in the next step, otherwise the system may go to q , rendering it unpredictable.
- As long as $I(\sigma) < 2^n - 1$, the controller must permit an action i such that $I(\sigma i) > I(\sigma)$. To see this, let $s_i \notin R(\sigma)$, then $R(\sigma i) = (R(\sigma) \cup \{s_i\}) \setminus \{s_0, \dots, s_{i-1}\}$. We shall assume that i is chosen minimally, so $I(\sigma i) = I(\sigma) + 1$.
- Therefore, after $2^n - 1$ steps, the controlled system will have performed a sequence $\hat{\sigma}$ with $I(\hat{\sigma}) = 2^n - 1$. The only possible course of action for the controller is to permit n from now on, i.e. $\rho = \hat{\sigma}n^\omega$. We then have $R(\hat{\sigma}n) = \{p, p'\}$, $R(\hat{\sigma}nn) = \{p', p''\}$, and $R(\hat{\sigma}nnn) = \{p''\}$.

Going backwards, we can now associate predictions with each prefix of ρ : $pred(\hat{\sigma}n^k) = \perp$ for $k \geq 3$, $pred(\hat{\sigma}nn) = \langle ?, 1 \rangle$, $pred(\hat{\sigma}n) = \langle 0, 2 \rangle$, and $pred(\sigma) = \langle 2^n - |\sigma|, 2 \rangle$ for every prefix σ of $\hat{\sigma}$. Thus, \mathcal{A}_n is 2^n -2-active predictable. Notice that the system could be made 2^n -1-active predictable if states s_0, \dots, s_{n-1} transitioned with n to p' instead, which we avoided simply to keep the drawing of the automaton planar.

Constant-size alphabet To see that the proof with a variable-size alphabet can be adapted to an alphabet of constant size, consider the LTS \mathcal{A}'_n in Figure 7. \mathcal{A}'_n has $\mathcal{O}(n)$ states and three observable and controllable actions $0, 1, a$ and two unobservable actions u and f . Initially, the LTS performs an a going to either p or r . The LTS then simulates \mathcal{A}_n of Figure 6, using a unary encoding, in the following sense: Let $code(i) = 1^i 0^{n-i} a$, for $i = 0, \dots, n$. The reader can verify, case-by-case, that for any two states $u, v \in \{p, p', s_0, \dots, s_{n-1}, q\}$ and $i \in \{0, \dots, n\}$, we have $u \xrightarrow{i} v$ in \mathcal{A}_n iff $u \xrightarrow{code(i)} v$ in \mathcal{A}'_n . Moreover, the controller must account for the possibility that the system has gone to state r . Then, to keep the controlled system live, the only possible sequences that the controller can enforce are $code(i)$ for $i = 0, \dots, n$, and we have $r \xrightarrow{code(i)} r$ for $i < n$. After the initial a , the controller must therefore admit $code(\hat{\sigma}n)$, for $\hat{\sigma}$ as in \mathcal{A}_n . On this basis, a closer look shows that \mathcal{A}'_n is k - l -active predictable for $k = 1 + (n + 1) \cdot 2^n$ and $l = n + 2$. \square

Note that Theorem 35 does not contradict Proposition 8, which establishes linear prediction bounds w.r.t. the number of states of \mathcal{A} . However, Proposition 8 talks about passive predictability, whereas Theorem 35 is about active predictability.

5 Conclusion and perspectives

We have extended the prediction paradigm by introducing parameters related to the number of observations before fault may or must occur. Within this framework, we have established that active predictability is EXPTIME-complete through a procedure for synthesising active predictors that builds a Büchi game. Solving this game is proved linear in the number of edges in the game. We have shown that if the observation threshold for *eventual* prediction is chosen large enough (namely $\geq 2^n$ with n the number of states in the system), then active predictability is equivalent to active safety. Furthermore we have exhibited a family of systems proving that this bound is tight.

Out of several possible extensions for the present results, three stand out as natural continuations. **First, we want to introduce a measure that quantifies the faultiness of the system, and then aim to find an active predictor that minimizes this criterium, or at least ensures a value below some threshold.** Second, we plan to study the notion of prediagnosis introduced in [2] that combines predictability and diagnosability for controllable systems. Finally, we also want to study active predictability for probabilistic systems, as we had previously done for diagnosis in [1].




Fig. 7: Variant of Figure 6 with constant-size alphabet, with $\Sigma_o = \Sigma_c = \{0, 1, a\}$.

References

1. N. Bertrand, E. Fabre, S. Haar, S. Haddad, and L. Hélouët. Active diagnosis for probabilistic systems. In *FOSSACS 2014, Grenoble, France*, volume 8412 of *LNCS*, pages 29–42, 2014.
2. N. Bertrand, S. Haddad, and E. Lefauchaux. Foundation of Diagnosis and Predictability in Probabilistic Systems. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'14)*, volume 29 of *LIPICs*, pages 417–429, New Delhi, India, December 2014.
3. D. Berwanger and L. Doyen. On the power of imperfect information. In *Proc. FSTTCS*, volume 2 of *LIPICs*, pages 73–82, Bangalore, India, 2008.
4. S. Böhm, S. Haar, S. Haddad, P. Hofman, and S. Schwoon. Active diagnosis with observable quiescence. In *Proc. CDC: 54th IEEE Conf. on Decision and Control*, pages 1663–1668, Osaka, Japan, December 2015.
5. L. Brandán Briones and A. Madalinski. Bounded predictability for faulty discrete event systems. In *30nd International Conference of the Chilean Computer Science Society, SCCC*, pages 142–146, Curico, Chile, November 2011.

6. C. G. Cassandras and S. Lafortune. *Introduction to Discrete Event Systems - Second Edition*. Springer, 2008.
7. F. Cassez and S. Tripakis. Fault diagnosis with static and dynamic observers. *Fundamenta Informaticae*, 88:497–540, 2008.
8. F. Cassez and S. Tripakis. Fault diagnosis with static and dynamic observers. *Fundam. Informaticae*, 88(4):497–540, 2008.
9. E. Chantry and Y. Pencolé. Monitoring and active diagnosis for discrete-event systems. In *Proc. SafeProcess'09*, pages 1545–1550, 2009.
10. E. Dallal and S. Lafortune. On most permissive observers in dynamic sensor activation problems. *IEEE Trans. Autom. Control.*, 59(4):966–981, 2014.
11. S. Genc and S. Lafortune. Predictability of event occurrences in partially-observed discrete-event systems. *Autom.*, 45(2):301–311, 2009.
12. E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.
13. S. Haar, S. Haddad, T. Melliti, and S. Schwon. Optimal constructions for active diagnosis. *Journal of Computer and System Sciences*, 83(1):101–120, 2017.
14. A. Madalinski and V. Khomenko. Predictability verification with parallel LTL-X model checking based on Petri net unfoldings. *IFAC Proceedings Volumes*, 45(20):1232 – 1237, 2012. 8th IFAC Symposium on Fault Detection, Supervision and Safety of Technical Processes.
15. M. Sampath, S. Lafortune, and D. Teneketzis. Active diagnosis of discrete-event systems. *IEEE Transactions on Automatic Control*, 43(7):908–929, July 1998.
16. M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, and D. Teneketzis. Diagnosability of discrete-event systems. *IEEE Trans. Aut. Cont.*, 40(9):1555–1575, 1995.
17. L. Ye, P. Dague, and F. Nouioua. Predictability Analysis of Distributed Discrete Event Systems. In *52nd IEEE Conference on Decision and Control*, pages 5009–5015, Florence, Italy, December 2013.
18. X. Yin and S. Lafortune. A uniform approach for synthesizing property-enforcing supervisors for partially-observed discrete-event systems. *IEEE Trans. Autom. Control.*, 61(8):2140–2154, 2016.
19. X. Yin and S. Lafortune. A general approach for optimizing dynamic sensor activation for discrete event systems. *Autom.*, 105:376–383, 2019.
20. X. Yin and Z. Li. Decentralized fault prognosis of discrete event systems with guaranteed performance bound. *Autom.*, 69:375–379, 2016.
21. T-S. Yoo and S. Lafortune. Polynomial-time verification of diagnosability of partially observed discrete-event systems. *IEEE Trans. Automat. Contr.*, 47(9):1491–1495, 2002.